

IT IS CLAIMED:

1. A method of switching packets to a server among a group of servers, wherein individual packets carry content associated with an application, said method comprising:

receiving a packet;

5 parsing application content from said packet;

maintaining a set of server load metrics associated with the servers among said group of servers;

selecting a destination server for said packet from among said group of servers, said selection being dependent on said set of server load metrics;

10 assigning to said packet a priority dependent on said parsed application content; and

transmitting said packet to the destination server according to a transmitting schedule dependent on said priority.

15 2. A method of switching packets to a server among a group of servers as in claim 1, further comprising:

determining a group of eligible servers among said group of servers, said determination being dependent on said parsed application content, and wherein

said destination server selection is restricted to said group of eligible servers.

20

3. A method of switching packets to a server among a group of servers, as in claim 1, further comprising:

determining an estimated application load on the destination server, said estimated application load being dependent on said parsed application content; and wherein

said selecting a destination server is also dependent on said estimated application load.

4. A method of switching packets to a server among a group of servers as in of
5 claim 1, wherein said transmitting schedule is also dependent on said set of server load metrics.

5. A method of switching packets to a server among a group of servers as in
claim1, wherein said application content includes that from a header for Hypertext
10 Transfer Protocol.

6. A method of switching packets to a server among a group of servers, wherein
individual packets carry content associated with an application, said method comprising:

- receiving a packet;
15 parsing application content from said packet;
maintaining a set of server load metrics associated with the servers among said group of servers;
determining a group of eligible servers among said group of servers, said determination being dependent on said parsed application content;
20 determining an estimated application load on the destination server, said estimated application load being dependent on said parsed application content;

selecting a destination server for said packet from among said group of eligible servers, said selection being dependent on said set of server load metrics and said estimated application load;

assigning to said packet a priority dependent on said parsed application content;
and

transmitting said packet to the destination server according to a transmitting
schedule dependent on said priority and said set of server load metrics.

5

7. A method of switching packets to a server among a group of servers as in
claim 6, wherein said application content includes that from a packet header for
Hypertext Transfer Protocol.

10 8. A method of switching packets to a server among a group of servers as in
anyone of claims 1-7, wherein:

the priority assigned to said packet is one of a first priority or a second priority
lower than said first priority; and

15 said transmitting schedule is such that whenever the destination server has a
workload above a first predetermined level, a first priority packet is transmitted thereto
without delay while a second priority packet is held back.

9. A method of switching packets to a server among a group of servers as in
claim 8, wherein:

20 said transmitting schedule is such that whenever the destination server has a
workload below said first predetermined level, a first priority packet or a second priority
packet is transmitted thereto without delay.

25 10. A method of switching packets to a server among a group of servers as in
anyone of claims 1-7, wherein:

the priority assigned to said packet is one of a hierarchy of priorities ranging from a highest priority to a lowest priority; and said method further comprising:

providing a set of predetermined workload levels for the destination server, said set of predetermined workload levels ranging from a highest level corresponding to the
5 highest priority to a lowest level corresponding to the lowest priority; and wherein

said transmitting schedule is such that whenever the destination server has a workload between a first of said levels and a second of said levels less than said first of said levels, a packet assigned the priority corresponding to said first of said levels is transmitted to the destination server without delay while a packet assigned the priority
10 corresponding to said second of said levels is held back.

11. A method of switching packets to a server among a group of servers as in claim 10, wherein:

the priority assigned to said packet includes one of zero priority; and

15 said transmitting schedule is such that a packet with zero priority is never transmitted to the destination server.

12. A method of switching packets to a server among a group of servers, wherein individual packets carry content associated with an application, said method comprising:

20 receiving a packet;

parsing application content from said packet;

assigning to said packet a priority dependent on said parsed application content;
and

dropping the packet if its assigned priority corresponds to a predetermined type.

25

13. A method of switching packets to a server among a group of servers, wherein individual packets carry content associated with an application, said method comprising:

receiving a packet;

parsing application content from said packet; and

- 5 enforcing security by reference to an access control list dependent on said parsed application content.

14. A method of switching packets to a server among a group of servers, wherein individual packets carry content associated with an application, as in claim 13, wherein
10 said enforcing of security includes dropping the packet if its parsed application content is of a first predetermined type.

15. A method of switching packets to a server among a group of servers, wherein individual packets carry content associated with an application, as in claim 13, wherein
15 said enforcing of security includes redirecting the packet to a predetermined location if its parsed application content is of a second predetermined type.

16. A method of switching packets to a server among a group of servers as in claim 12, wherein said application content includes that from a packet header for
20 Hypertext Transfer Protocol.

17. A switch for switching packets to a server among a group of servers, wherein individual packets carry content associated with an application, comprising:

a buffer memory for storing individual packets;

5 a content class table defining a set of content classes with respect to a corresponding set of application content;

a class policy table associating a priority to a given content class;

a server table providing a set of server load metrics associated with the servers among said group of servers;

10 a packet classifier for assigning a content class to a packet, said packet classifier parses the application content from said packet and determines said content class by reference to said content class table for said parsed application content;

a load balancer for selecting a destination server from among said group of servers based on a set of server load metrics by reference to said server table for said group of servers;

15 a Quality-of-Service controller for scheduling transmission of said packet to the destination server according to a priority assigned to said packet by reference to said class policy table for said content class; and

20 a transmitter, responsive to said load balancer and said Quality-of-Service controller for releasing said packet from said buffer to the destination server in the order the packet is scheduled.

18. A switch for switching packets to a server among a group of servers as in claim 17, further comprising:

25 a server group table associating a content class to a group of eligible servers and wherein

said load balancer selects a destination server from among said group of eligible servers.

19. A switch for switching packets to a server among a group of servers as in
5 claim 17, further comprising:

an application load table associating a content class to an estimated application load on the selected server; and wherein:

the selected server has a server load metric that is dependent on said estimated application load.
10

20. A switch for switching packets to a server among a group of servers as in claim 17, wherein:

said transmission schedule of said packet to the destination server is also based on said set of server load metrics.

15

21. A switch for switching packets to a server among a group of servers as in claim 17, wherein said application content includes that from a packet header for Hypertext Transfer Protocol.

22. A switch for switching a packet to a server among a group of servers, wherein individual packets carry content associated with an application, comprising:

a buffer memory for storing said packet;

a content class table defining a set of content classes with respect to a corresponding set of application content;

a class policy table associating a priority to a given content class;

a server group table associating a content class to a group of eligible servers;

an application load table associating a content class to an estimated application load on the selected server;

- 5 a server table providing a set of server load metrics associated with the servers among said group of servers;

a packet classifier for assigning a content class to said packet, said packet classifier parses the application content from said packet and determines said content class by reference to said content class table for said parsed application content;

- 10 a load balancer for selecting a destination server from among said group of servers, said selecting being restricted to a group of eligible servers by reference to said server group table for said content class, and said selecting being based on said set of server load metrics by reference to said server table for said group of eligible servers and based on an estimated application load by reference to said application load table for said
15 content class;

a Quality-of-Service controller for scheduling transmission of said packet to the destination server according to a priority assigned to said packet by reference to said class policy table for said content class and said set of server load metrics; and

- a transmitter, responsive to said load balancer and said Quality-of-Service
20 controller for releasing said packet from said buffer to the destination server in the order the packet is scheduled.

23. A switch for switching packets to a server among a group of servers as in claim 22, wherein said application content includes that from a packet header for
25 Hypertext Transfer Protocol.

24. A switch for switching packets to a server among a group of servers as in claims 17-22, wherein

the priority assigned to said packet is one of a first priority or a second priority lower than said first priority; and

- 5 said transmitting schedule is such that whenever the destination server has a workload above a first predetermined level, a first priority packet is transmitted thereto without delay while a second priority packet is held back.

25. A switch for switching packets to a server among a group of servers as in
10 claim 24, wherein:

 said transmitting schedule is such that whenever the destination server has a workload below said first predetermined level, a first priority packet or a second priority packet is transmitted thereto without delay.

- 15 26. A switch for switching packets to a server among a group of servers as in anyone of claims 17-22, wherein:

 the priority assigned to said packet is one of a hierarchy of priorities ranging from a highest priority to a lowest priority; and said switch further comprising:

- 20 a set of predetermined workload levels for the destination server, said set of predetermined workload levels ranging from a highest level corresponding to the highest priority to a lowest level corresponding to the lowest priority; and wherein

- said transmitting schedule is such that whenever the destination server has a workload between a first of said levels and a second of said levels less than said first of said levels, a packet assigned the priority corresponding to said first of said levels is
25 transmitted to the destination server without delay while a packet assigned the priority corresponding to said second of said levels is held back.

27. A switch for switching packets to a server among a group of servers as in claim 26, wherein:

the priority assigned to said packet includes one of zero priority; and

5 said transmitting schedule is such that a packet with zero priority is never transmitted to the destination server.

28. A switch for switching packets to a server among a group of servers, wherein individual packets carry content associated with an application, comprising:

10 a buffer memory for storing individual packets;

a content class table defining a set of content classes with respect to a corresponding set of application content;

a class policy table associating a priority to a given content class;

15 a packet classifier for assigning a content class to a packet, said packet classifier parses the application content from said packet and determines said content class by reference to said content class table for said parsed application content; and

a transmit controller for dropping the packet if its assigned priority corresponds to a predetermined type.

20 29. A switch for switching packets to a server among a group of servers, wherein individual packets carry content associated with an application, comprising:

a buffer memory for storing individual packets;

a content class table defining a set of content classes with respect to a corresponding set of application content;

an access control list table associating security access instructions to a given content class;

a packet classifier for assigning a content class to a packet, said packet classifier parses the application content from said packet and determines said content class by reference to said content class table for said parsed application content; and

a transmit controller responsive to said access control list table for enforcing securing access of the packet.

30. A switch for switching packets to a server among a group of servers as in claim 29, wherein said transmit controller drops the packet when its parsed application content is of a first content class.

31. A switch for switching packets to a server among a group of servers as in claim 29, wherein said transmit controller redirects the packet to a predetermined location when its parsed application content is of a second content class.

32. A switch for switching packets to a server among a group of servers as in claim 28, wherein said application content includes that from a packet header for Hypertext Transfer Protocol.

33. A switch for switching packets to a server among a group of servers, wherein individual packets carry content associated with an application, comprising:

means for receiving a packet;

means for parsing application content from said packet;

means for maintaining a set of server load metrics associated with the servers among said group of servers;

means for selecting a destination server for said packet from among said group of servers, said selection being dependent on said set of server load metrics;

5 means for assigning to said packet a priority dependent on said parsed application content; and

means for transmitting said packet to the destination server according to a transmitting schedule dependent on said priority.

10 34. A switch for switching packets to a server among a group of servers as in claim 33, further comprising:

means for determining a group of eligible servers among said group of servers, said determination being dependent on said parsed application content, and wherein

said destination server selection is restricted to said group of eligible servers.

15

35. A switch for switching packets to a server among a group of servers, as in claim 33, further comprising:

means for determining an estimated application load on the destination server, said estimated application load being dependent on said parsed application content; and

20 wherein

said selecting a destination server is also dependent on said estimated application load.

36. A switch for switching packets to a server among a group of servers as in of claim 33, wherein said transmitting schedule is also dependent on said set of server load metrics.

5 37. A switch for switching packets to a server among a group of servers as in anyone of claim 33, wherein said application content includes that from a header for Hypertext Transfer Protocol.

10 38. A switch for switching packets to a server among a group of servers, wherein individual packets carry content associated with an application, comprising:

means for receiving a packet;

means for parsing application content from said packet;

means for maintaining a set of server load metrics associated with the servers among said group of servers;

15 means for determining a group of eligible servers among said group of servers, said determination being dependent on said parsed application content;

means for determining an estimated application load on the destination server, said estimated application load being dependent on said parsed application content;

20 means for selecting a destination server for said packet from among said group of eligible servers, said selection being dependent on said set of server load metrics and said estimated application load;

means for assigning to said packet a priority dependent on said parsed application content; and

25 means for transmitting said packet to the destination server according to a transmitting schedule dependent on said priority and said set of server load metrics.

39. A switch for switching packets to a server among a group of servers as in claim 38, wherein said application content includes that from a packet header for Hypertext Transfer Protocol.

5

40. A switch for switching packets to a server among a group of servers as in anyone of claims 33-39, wherein:

the priority assigned to said packet is one of a first priority or a second priority lower than said first priority; and

10 said transmitting schedule is such that whenever the destination server has a workload above a first predetermined level, a first priority packet is transmitted thereto without delay while a second priority packet is held back.

41. A switch for switching packets to a server among a group of servers as in
15 claim 40, wherein:

said transmitting schedule is such that whenever the destination server has a workload below said first predetermined level, a first priority packet or a second priority packet is transmitted thereto without delay.

20 42. A switch for switching packets to a server among a group of servers as in anyone of claims 33-39, wherein:

the priority assigned to said packet is one of a hierarchy of priorities ranging from a highest priority to a lowest priority; and said method further comprising:

means for providing a set of predetermined workload levels for the destination
25 server, said set of predetermined workload levels ranging from a highest level

corresponding to the highest priority to a lowest level corresponding to the lowest priority; and wherein

said transmitting schedule is such that whenever the destination server has a workload between a first of said levels and a second of said levels less than said first of said levels, a packet assigned the priority corresponding to said first of said levels is transmitted to the destination server without delay while a packet assigned the priority corresponding to said second of said levels is held back.

43. A switch for switching packets to a server among a group of servers as in claim 42, wherein:

the priority assigned to said packet includes one of zero priority; and

said transmitting schedule is such that a packet with zero priority is never transmitted to the destination server.

44. A switch for switching packets to a server among a group of servers, wherein individual packets carry content associated with an application, said method comprising:

means for receiving a packet;

means for parsing application content from said packet;

means for assigning to said packet a priority dependent on said parsed application content; and

means for dropping the packet if its assigned priority corresponds to a predetermined type.

45. A switch for switching packets to a server among a group of servers, wherein individual packets carry content associated with an application, comprising:

means for receiving a packet;

means for parsing application content from said packet; and

means for providing application layer security by enforcing access control dependent on said parsed application content.

5

46. A switch for switching packets to a server among a group of servers, wherein individual packets carry content associated with an application, as in claim 45, wherein:

said means for providing application layer security includes dropping the packet if its parsed application content is of a first predetermined type.

10

47. A switch for switching packets to a server among a group of servers, wherein individual packets carry content associated with an application, as in claim 45, wherein:

said means for providing application layer security includes redirecting the packet to a predetermined location if its parsed application content is of a second predetermined type.

15

48. A switch for switching packets to a server among a group of servers as in claim 44, wherein said application content includes that from a packet header for Hypertext Transfer Protocol.

20

49. A method of switching packets to a server among a group of servers including a newly added server participating in load balancing, comprising:

maintaining a set of server load metrics associated with the servers among said group of servers, wherein a load metric associated with a server is dependent on the number of connections currently made hereto, and a disparity initially exists in server

25

load metrics between said newly added server and that of other servers among said group of servers;

- 5 multiplying a time-varying factor to the server load metric for said newly added server such that said disparity is substantially reduced to below a predetermined level over a predefined period of time;

selecting for a packet a destination server having a minimum server load metric among said group of servers; and

transmitting said packet to the destination server.

- 10 50. A method of switching packets as in claim 49, wherein:

said time-varying factor decreases from a first value substantially above unity to unity over said predefined period of time.

51. A method of switching packets as in claim 50, wherein:

- 15 said first value is given by 2^k , where k is an integer, and is reduced by a factor of two at a time until said first value becomes unity.

52. A method of switching packets to a server among a group of servers including a newly added server participating in load balancing, comprising:

- 20 maintaining a set of server load metrics associated with the servers among said group of servers, wherein a load metric associated with a server is dependent on the number of connections currently made hereto, and a disparity initially exists in server load metrics between said newly added server and that of other servers among said group of servers;

multiplying a factor of 2^k , with k being an integer, to the server load metric for said newly added server;

reducing said factor of 2^k by a factor of two at predefined periods until unity is reached;

- 5 selecting for said packet a destination server having a minimum server load metric among said group of servers; and

transmitting said packet to the destination server.

53. A method of switching packets as in anyone of claims 49-52, wherein:

- 10 a server weight is associated with each server to indicate a capacity capability of each server; and

the load metric of each server is weighted by its associated server weight.

54. A method of switching packets as in anyone of claims 49-52, wherein
15 individual packets carry content associated with an application, said method further comprising:

receiving a packet;

parsing application content from said packet;

- determining a group of eligible servers among said group of servers, said
20 determination being dependent on said parsed application content, and wherein

said destination server selection is restricted to said group of eligible servers.

55. A method of switching packets as in anyone of claims 49-52, further comprising:

determining an estimated application load on the destination server, said estimated application load being dependent on said parsed application content; and wherein

the server load metric for the destination server is also dependent on said estimated application load.

5

56. A method of switching packets as in claim 55, wherein said application content includes that from a header for Hypertext Transfer Protocol.

57. A method of switching packets as in anyone of claims 49-52, said method
10 comprising:

receiving a packet;

parsing application content from said packet;

determining a group of eligible servers among said group of servers, said
determination being dependent on said parsed application content;

determining an estimated application load on the destination server, said estimated
15 application load being dependent on said parsed application content; and wherein

the server load metric for the destination server is also dependent on said
estimated application load.

58. A method of switching as in claim 57, wherein said application content
20 includes that from a packet header for Hypertext Transfer Protocol.

59. A method of switching packets to a server among a group of servers, wherein
individual packets carry content associated with an application, said method comprising:

maintaining a set of server load metrics associated with the servers among said group of servers,

receiving a packet during a TCP session;

parsing application content from said packet;

- 5 updating said set of server load metrics only when the TCP session is not in an initial handshaking phase, said updating incorporating an estimated application load determined with reference to the application content parsed from said packet;

selecting a destination server for said packet from among said group of servers, said selection being dependent on said set of server load metrics;

- 10 assigning to said packet a priority dependent on application content parsed from said packet only when the TCP session is not in an initial handshaking phase, otherwise assigning to said packet a default priority; and

transmitting said packet to the destination server according to a transmitting schedule dependent on said priority

15

60. A method of switching packets to a server among a group of servers as in claim 59, wherein said default priority is a highest among a group of possible priority types.

- 20 61. A method of switching packets to a server among a group of servers as in claim 59, further comprising:

determining a group of eligible servers among said group of servers, said determination being dependent on said parsed application content, and wherein

said destination server selection is restricted to said group of eligible servers.

25

62. A method of switching packets to a server among a group of servers, as in claim 59, further comprising:

determining an estimated application load on the destination server, said estimated application load being dependent on said parsed application content; and wherein

5 said selecting a destination server is also dependent on said estimated application load.

63. A method of switching packets to a server among a group of servers as in of claim 59, wherein said transmitting schedule is also dependent on said set of server load
10 metrics.

64. A method of switching packets to a server among a group of servers as in claims 59, wherein said application content includes that from a header for Hypertext Transfer Protocol.
15

65. A method of switching packets to a server among a group of servers as in claims 60, wherein said application content includes that from a header for Hypertext Transfer Protocol.

20 66. A method of switching packets to a server among a group of servers as in claims 61, wherein said application content includes that from a header for Hypertext Transfer Protocol.

67. A method of switching packets to a server among a group of servers as in claims 59, wherein said application content includes that from a header for Hypertext Transfer Protocol.

5 68. A method of switching packets to a server among a group of servers as in claims 63, wherein said application content includes that from a header for Hypertext Transfer Protocol.

10 69. A method of switching packets to a server among a group of servers as in claim 60, further comprising:

determining a group of eligible servers among said group of servers, said determination being dependent on said parsed application content, and wherein

said destination server selection is restricted to said group of eligible servers.

15 70. A method of switching packets to a server among a group of servers, as in claim 60, further comprising:

determining an estimated application load on the destination server, said estimated application load being dependent on said parsed application content; and wherein

20 said selecting a destination server is also dependent on said estimated application load.

71. A method of switching packets to a server among a group of servers as in of claim 60, wherein said transmitting schedule is also dependent on said set of server load metrics.

25

72. A method of switching packets to a server among a group of servers as in claim 69, wherein said application content includes that from a header for Hypertext Transfer Protocol.

- 5 73. A method of switching packets to a server among a group of servers as in claim 70, wherein said application content includes that from a header for Hypertext Transfer Protocol.

- 10 74. A method of switching packets to a server among a group of servers as in claim 71, wherein said application content includes that from a header for Hypertext Transfer Protocol.